arXiv:1602.07640v1 [stat.CO] 24 Feb 2016

A Variational Algorithm for Bayesian Variable Selection

Xichen Huang XHUANG43@ILLINOIS.EDU
Department of Statistics
University of Illinois at Urbana-Champaign
Champaign, IL 61820, USA

Jin Wang JINWANG8@ILLINOIS.EDU
Department of Statistics
University of Illinois at Urbana-Champaign
Champaign, IL 61820, USA

Feng Liang LIANGF@ILLINOIS.EDU
Department of Statistics at Urbana-Champaign
University of Illinois
Champaign, IL 61820, USA

Editor:

Abstract

There has been an intense development on the estimation of a sparse regression coefficient vector in statistics, machine learning and related fields. In this paper, we focus on the Bayesian approach to this problem, where sparsity is incorporated by the so-called spike-and-slab prior on the coefficients. Instead of replying on MCMC for posterior inference, we propose a fast and scalable algorithm based on variational approximation to the posterior distribution. The updating scheme employed by our algorithm is different from the one proposed by Carbonetto and Stephens (2012). Those changes seem crucial for us to show that our algorithm can achieve asymptotic consistency even when the feature dimension diverges exponentially fast with the sample size. Empirical results have demonstrated the effectiveness and efficiency of the proposed algorithm.

Keywords: variable selection, variational approximation, spike-and-slab prior, consistency, Bayesian consistency

1. Introduction

Consider a standard linear regression problem, where we model Y, a continuous response variable, by a linear function of a set of p features $(X_1, \ldots, X_p)$ via

$$Y = X_1\beta_1 + \cdots X_p\beta_p + \epsilon.$$

In the past three decades or so, there has been an intense development on the estimation of a sparse regression model. Here "sparse" means that only a small fraction of β_j's is believed to be non-zero. Identifying the set $S = \{j : \beta_j \neq 0, j = 1, \ldots, p\}$ is often referred to as the variable selection problem.

The current approaches to variable selection can be roughly divided into two categories. One category contains approaches based on penalized likelihood, including the classical variable selection procedures like AIC/BIC and the more recent ones like LASSO (Tibshirani, 1994) and SCAD (Fan and Li, 2001). As the name suggested, the penalized likelihood approach estimates the regression parameter by minimizing the log-likelihood plus some penalty function on $\boldsymbol{\beta}$. With a proper choice of the penalty function, the solution $\hat{\boldsymbol{\beta}}$ will have some of its components to be exactly zero, that is, parameter estimation and variable selection are carried out simultaneously. For an overview of the recent developments on penalized likelihood approaches to variable selection in high dimensions, see Fan and Lv (2010) and Bühlmann and van de Geer (2011).

We focus on the other category, the Bayesian approach, which starts with a hierarchical prior on all the unknown parameters. For example, a widely used prior on $\boldsymbol{\beta}$ is the so-called spike-and-slab prior (Mitchell and Beauchamp, 1988):

$$\beta_j|\gamma_j, \sigma_2 \sim \gamma_j \, \mathrm{N}(0, v_1\sigma^2) + (1 - \gamma_j)\delta_0, \quad j = 1, \ldots, p, \tag{1}$$

where δ_0 denotes a point mass at 0, and $\gamma_j = 1$ if the j-th variable is included and 0 otherwise. The p-dimensional binary vector $\boldsymbol{\gamma}$, which serves as a model index for all the 2^p sub-models, is then modeled by a product of *i.i.d.* Bernoulli distributions with parameter θ.

An advantage of the Bayesian approach is that, in addition to the posterior distribution on $\boldsymbol{\beta}$, we can also obtain a posterior distribution on all the sub-models. For example, we can discuss the probability of a sub-model $\boldsymbol{\gamma}$ or the inclusion probability of a particular feature, which can be of more interest than a point estimate of $\boldsymbol{\beta}$. Further, for prediction, it is well-known that model combination or aggregation has a better performance than a single model (Breiman, 2001). The Bayesian approach for variable selection gives rise to a natural averaging scheme: the prediction from various sub-models can be averaged with respect to their posterior probabilities (Raftery et al., 1998; Clyde and George, 2004).

Despite the aforementioned advantages, in practice, Bayesian variable selection is less preferable than those penalization algorithms. A major disadvantage of Bayesian variable selection is the computing cost. The posterior distribution usually does not have a closed-form expression, so posterior inference has to reply on MCMC, which could be time consuming especially when the number of predictors is large.

In this paper, we propose a variational algorithm for Bayesian variable selection. It is a deterministic algorithm, seeking an approximation of the true posterior distribution over $(\boldsymbol{\beta}, \boldsymbol{\gamma})$, instead of running an MCMC chain. It converges very fast and can scale for large sized data sets. Our work is motivated by an earlier variational algorithm proposed by Carbonetto and Stephens (2012). The two algorithms have the same prior specification and the same set of variational parameters Θ. The two algorithms, however, update the variational parameters differently. In the algorithm by Carbonetto and Stephens (2012), the parameters associated with each feature are updated sequentially given the others; such a *component-wise updating* scheme is prone to error accumulation especially when p is large and predictors are correlated. In our algorithm, all features are updated simultaneously, which we refer to as the *batch-wise updating* scheme, therefore is more robust to errors and correlations among predictors. Indeed, the batch-wise updating scheme employed by our algorithm turns out to be crucial for us to show our algorithm achieves both frequentist consistency and Bayesian consistency even when p diverges at an exponential rate of the sample size n. To the best of our knowledge, no asymptotic results on variational algorithms for Bayesian variable selection are available in the literature.

The remaining of the paper is arranged as follows: Section 2 presents the two variational Bayes (VB) algorithms; Section 3 investigates the asymptotic properties of our new algorithm; Empirical results are given in Section 4 and Section 5, and conclusions are given in Section 6.

1.1 Notation.

We define some symbols that will be used in the following sections. For sequences $\{a_n\}_{n=1}^{\infty}$ and $\{b_n\}_{n=1}^{\infty}$, we write

- $a_n = O(b_n)$, if $\exists\ c \in \mathbb{R}^+$ and $n_0 \in \mathbb{N}$, s.t. $|a_n/b_n| \leq c$, $\forall n \geq n_0$;
- $a_n = o(b_n)$, if $\lim_{n\to\infty} a_n/b_n = 0$;
- $a_n \asymp b_n$, if $\exists\ c_1, c_2 \in \mathbb{R}^+$ and $n_0 \in \mathbb{N}$, s.t. $c_1 \leq |a_n/b_n| \leq c_2$, $\forall n \geq n_0$;
- $a_n \prec b_n$ if $a_n = o(b_n)$, and $a_n \preceq b_n$ if $a_n = O(b_n)$.

For a random variable sequence $\{X_n\}_{n=1}^{\infty}$ and a constant sequence $\{a_n\}_{n=1}^{\infty}$, we write $X_n = O_P(a_n)$ if $\forall \varepsilon > 0$, $\exists M > 0$ s.t. $P\left(|\frac{X_n}{a_n}| > M\right) < \varepsilon$, $\forall n$, and $X_n = o_p(a_n)$ when $\lim_{n\to\infty} P\left(|\frac{X_n}{a_n}| \geq \varepsilon\right) = 0$, $\forall \varepsilon > 0$. For $a, b \in \mathbb{R}$, we write $a \bigvee b$ to represent the larger number of a and b, and $a \bigwedge b$ to represent the smaller one of a and b.

2. Variational Approximation

2.1 The Model

Represent the linear regression model in a matrix form:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{2}$$

where $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_n)^T$ is a vector that contains n *i.i.d.* random errors generated from a normal distribution $\mathrm{N}(0, \sigma^2)$, $\mathbf{y}$ is the response vector of length n, $\mathbf{X} = (x_{ij})$ is an $n \times p$ design matrix, and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)^T$ is the coefficient vector of length p. Like in many other variable selection algorithms, we center and scale the data as follows:

$$\sum_i y_i = 0, \quad \sum_i x_{ij} = 0, \quad \sum_i x_{ij}^2 = \|\mathbf{X}_j\|_2^2 = n,$$

where $\mathbf{X}_j$ denotes the j-th column of $\mathbf{X}$.

The hierarchical prior is specified as follows:

$$\begin{aligned} \beta_j|\gamma_j &\sim \gamma_j \,\mathrm{N}(0, v_1\sigma^2) + (1-\gamma_j)\delta_0, \\ \gamma_j &\overset{i.i.d.}{\sim} \mathrm{Bern}(\theta), \\ \sigma^2 &\sim \mathrm{IG}\left(\frac{\nu}{2}, \frac{\nu\lambda}{2}\right), \\ \theta &\sim \mathrm{Beta}(a_0, b_0), \end{aligned} \tag{3}$$

where $j = 1, \ldots, p$, and ν, λ, a_0 and b_0 are hyper-parameters.

2.2 A Variational EM Algorithm

Variational methods have been widely used in different models, such as the Graphic models (Jordan et al., 1999). In the ordinary variational Bayesian approach (Bishop, 2006), an approximating distribution Q of all the latent variables and parameters, which takes a factorized form of $\prod_j Q_j$, is selected from a restricted family of distributions $\mathcal{Q}$, such that the negative KL-divergence from the true posterior P to Q is maximized, i.e.,

$$\max_{Q \in \mathcal{Q}} \mathbb{E}^Q \log \frac{P}{Q} = \int Q \log \frac{P}{Q} \,\mathrm{d}Q.$$

Then one can solve each Q_j sequentially by fixing other Q's until convergence.

Our variational algorithm is a hybrid of Expectation-Maximization (EM) and variational, same as the one used by Blei et al. (2003) for topic models. Next we give a general description of the framework we use for posterior inference.

Let (Θ_1, Θ_2) denote the set of parameters of interest, and η denote the hyper-parameters. The goal is to obtain an approximation of the posterior distribution on (Θ_1, Θ_2). Define the following objective function

$$\Omega(q_1, q_2, \eta) = \mathbb{E}^{q_1,q_2}_{\Theta_1,\Theta_2} \log \frac{\pi(\Theta_1, \Theta_2, \eta|\text{Data})}{q_1(\Theta_1)q_2(\Theta_2)}, \tag{4}$$

where q_1 and q_2 are distributions on Θ_1 and Θ_2 respectively. Our goal is to find q_1, q_2, and a point estimate $\hat{\eta}$ to maximize the objective function. We will refer to the estimate $\hat{\eta}$ as the Maximum a posteriori (MAP) estimate: if we optimize (4) with respect to $q(\Theta_1, \Theta_2)$ and η, instead of restricting $q(\Theta_1, \Theta_2)$ to take a product form, then the corresponding $\hat{\eta}$ will be exactly the MAP estimate of η.

Applying the framework above on the Bayesian variable selection model, we estimate the MAP for σ^2 and θ, and approximate the posteriors for β_j's and γ_j's. The approximating posterior distribution of $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ takes the following form:

$$q(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \prod_{j=1}^{p} q_j(\beta_j, \gamma_j) = \prod_{j=1}^{p} [\phi_j f_j(\beta_j)]^{\gamma_j} [(1-\phi_j)\delta_0(\beta_j)]^{1-\gamma_j},$$

where $f_j(\beta_j)$ is a probability density function. That is, we approximate the posterior distribution of β_j by $\beta_j = 0$ with probability $1-\phi_j$, and $\beta_j \neq 0$ following a continuous distribution with probability ϕ_j.

Given all the information above, we define the following objective function for this problem

$$\Omega(q_1, \ldots, q_p, \theta, \sigma^2) = \mathbb{E}^{q_1,\ldots,q_p} \log \frac{p(\mathbf{y}|\boldsymbol{\beta}, \sigma^2)p(\boldsymbol{\beta}|\boldsymbol{\gamma})p(\boldsymbol{\gamma}|\theta)\pi(\theta)\pi(\sigma^2)}{\prod_{j=1}^{p} [\phi_j f_j(\beta_j)]^{\gamma_j} [(1-\phi_j)\delta_0(\beta_j)]^{1-\gamma_j}}.$$

2.3 Algorithm 1 : Component-wise VB

The first algorithm is similar to the variational algorithm proposed by Carbonetto and Stephens (2012). In detail, we iteratively update the approximating distributions of $q_j(\beta_j, \gamma_j)$'s, and the MAP estimates $\hat{\theta}$ and $\hat{\sigma}^2$. Since the algorithm loops over the p dimensions feature by feature, we refer to it as a "component-wise" VB algorithm, to highlight its difference with Algorithm 2, which we shall propose.

2.3.1 Updating Equations

Update $q_j(\beta_j, \gamma_j)$**.** For some $j \in \{1, \ldots, p\}$, by fixing other approximating distributions and point estimates, we maximize the objective function with respect to q_j. As shown in Carbonetto and Stephens (2012), $f_j(\beta_j)$ is the probability density function of a

Normal distribution $\mathrm{N}(\beta_j|\mu_j, \sigma_j^2)$ (albeit we do not assume f_j to be a normal at the beginning) with

$$\mu_j = \frac{\mathbf{X}_j^T \mathbb{E}_{[-j]}\left(\mathbf{y} - \sum_{l\neq j} \mathbf{X}_l \beta_l\right)}{n + \frac{1}{v_1}},$$
$$\sigma_j^2 = \frac{\hat{\sigma}^2}{n + \frac{1}{v_1}},$$

where $\mathbb{E}_{[-j]}$ denotes the expectations over all the β_l's with $l \neq j$ with respect to the variational distributions. By symmetry, we know that the other $f_j(\beta_j)$'s are also Normal density functions. As such, we can write μ_j as

$$\mu_j = \frac{\left(\mathbf{y} - \mathbf{X}_{[-j]}\bar{\boldsymbol{\beta}}_{[-j]}\right)^T \mathbf{X}_j}{n + \frac{1}{v_1}}, \tag{5}$$

where $\mathbf{X}_{[-j]}$ denotes the design matrix without the j-th column, and $\bar{\boldsymbol{\beta}} = (\phi_1\mu_1, \ldots, \phi_p\mu_p)^T$ is the mean of $\boldsymbol{\beta}$ w.r.t. $q(\boldsymbol{\beta}, \boldsymbol{\gamma})$.

The log-odds of ϕ_j can be updated as

$$\mathrm{Logit}(\phi_j) = \mathrm{Logit}(\hat{\theta}) + \frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2} + \frac{\mu_j^2}{2\sigma_j^2}.$$

Update $\hat{\theta}$. The point estimate of θ is updated by

$$\hat{\theta} = \frac{\sum_{j=1}^p \phi_j + a_0 - 1}{p + a_0 + b_0 - 2}. \tag{6}$$

Update $\hat{\sigma}^2$. The point estimate of σ^2 is updated by

$$\hat{\sigma}^2 = \frac{\|\mathbf{y} - \mathbf{X}\bar{\boldsymbol{\beta}}\|_2^2 + \sum_{j=1}^p[(n(1-\phi_j) + 1/v_1)\phi_j\mu_j^2 + (n + 1/v_1)\phi_j\sigma_j^2] + \nu\lambda}{n + \prod_{j=1}^p \phi_j + \nu + 2}. \tag{7}$$

2.3.2 The Drawback of Algorithm 1 in High Dimension

To reveal the potential drawback of Algorithm 1 when being applied on a high-dimensional data set, we examine its asymptotic property.

Assume the response $\mathbf{y}$ is generated from the normal linear regression model (2) with $\boldsymbol{\beta}^* \in \mathbb{R}^p$ being the true regression coefficients. Consider a relatively easy setting where the minimal eigenvalue of $\mathbf{X}^T\mathbf{X}$ is $O(n)$, i.e., the correlation among columns of $\mathbf{X}$ is small,

Algorithm 1 Component-wise VB

initialize $(\mu_1, \ldots, \mu_p)$, $(\sigma_1^2, \ldots, \sigma_p^2)$, $(\phi_1, \ldots, \phi_p)$, $\hat{\theta}$, and $\hat{\sigma}^2$.

repeat

for j in $1:p$ **do**

$\mu_j \leftarrow \frac{\left(\mathbf{y}-\sum_{l\neq j}\mathbf{X}_l\phi_l\mu_l\right)^T\mathbf{X}_j}{n+\frac{1}{v_1}}$

$\sigma_j^2 \leftarrow \frac{\hat{\sigma}^2}{n+\frac{1}{v_1}}$

$\phi_j \leftarrow \text{Logit}^{-1}\left\{\log\frac{\hat{\theta}}{1-\hat{\theta}} - \frac{1}{2}\log\frac{v_1\hat{\sigma}^2}{\sigma_j^2} + \frac{\mu_j^2}{2\sigma_j^2}\right\}$

end for

$\hat{\theta} \leftarrow \frac{\sum_{j=1}^p \phi_j + a_0 - 1}{p+a_0+b_0-2}$

$\hat{\sigma}^2 \leftarrow \frac{\|\mathbf{y}-\mathbf{X}\bar{\boldsymbol{\beta}}\|_2^2 + \sum_{j=1}^p[(n(1-\phi_j)+1/v_1)\phi_j\mu_j^2 + (n+1/v_1)\phi_j\sigma_j^2] + \nu\lambda}{n+\prod_{j=1}^p\phi_j+\nu+2}$

until Converge

return $(\mu_1, \ldots, \mu_p)$, $(\sigma_1^2, \ldots, \sigma_p^2)$, $(\phi_1, \ldots, \phi_p)$, $\hat{\theta}$, and $\hat{\sigma}^2$

and our starting values for μ_j's and ϕ_j's are very close to the truth: $\phi_j = 0.99$ if $\beta_j^* \neq 0$, $\phi_j = 0.01$ if $\beta_j^* = 0$, and

$$\mu_j - \beta_j^* = O_P\left(\frac{1}{\sqrt{n}}\right), \text{ for all } j = 1, \ldots, p.$$

Then, suppose we are updating the parameters associated with the j-th feature $(\mu_j, \sigma_j^2, \phi_j)$. After the update, will μ_j and ϕ_j still be close to the truth?

From Eq (5) we have

$$\begin{aligned}\mu_j =& \frac{\mathbf{X}_j^T(\mathbf{y} - \mathbf{X}_{[-j]}\bar{\boldsymbol{\beta}}_{[-j]})}{n + \frac{1}{v_1}} \\ =& \frac{nv_1}{nv_1+1}\Big[\beta_j^* + \frac{1}{n}\mathbf{X}_j^T\mathbf{X}_{[-j]}\big(\boldsymbol{\beta}_{[-j]}^* - \bar{\boldsymbol{\beta}}_{[-j]}\big) + \frac{1}{n}\mathbf{X}_j^T\boldsymbol{\epsilon}\Big].\end{aligned}$$

Suppose v_1 is chosen such that $v_1 n \to \infty$, a condition required for consistency as will be made clear in our analysis in Section 3. Then, we have

$$\mu_j = \beta_j^* + O_P\left(\frac{p}{\sqrt{n}}\right). \tag{8}$$

The result above shows the price we pay for Algorithm 1: even if we start with μ_j within a $1/\sqrt{n}$ ball around the truth β_j^*, after the update, the new μ_j could be very far away from β_j^* when p is large, due to the accumulation of the errors from other dimensions via $\mathbf{X}_{[-j]}\bar{\boldsymbol{\beta}}_{[-j]}$.

Next we examine how ϕ_j is affected by the update. Suppose the j-th feature is an irrelevant feature, i.e., $\beta_j^* = 0$. The new log-odds of ϕ_j is computed as

$$\begin{aligned}\text{Logit}(\phi_j) =& \text{Logit}(\hat{\theta}) + \frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2} + \frac{\mu_j^2}{2\sigma_j^2} \\ =& O(1) - \frac{1}{2}\log(v_1 n + 1) + \mu_j^2 \frac{n + \frac{1}{v_1}}{2\hat{\sigma}^2},\end{aligned}$$

where $\text{Logit}(\hat{\theta}) = O(1)$ as long as we do not start with $\hat{\theta} = 0$ or 1. Since $\mu_j^2 = O_P\left(p^2/n\right)$ by (8), we have

$$2\,\text{Logit}(\phi_j) = -\log(v_1 n + 1) + O_P(p^2). \tag{9}$$

When p is very large, the right hand side of (9) may be positive. That is, the new ϕ_j could be bigger than 0.5, although we start with $\phi_j = 0.01$, a value that is very close to the truth.

Our analysis above is not rigorous, but it clearly reveals an issue with Algorithm 1: the noise can accumulate due to the feature by feature updating scheme. To address this issue, we propose another algorithm which updates $(\mu_1, \ldots, \mu_p)$ simultaneously for all p features.

2.4 Algorithm 2: Batch-wise VB

Recall the variational parameters we need to update are $\{\mu_j, \phi_j, \sigma_j^2\}_{j=1}^p$. At iteration t, instead of updating the triplet $\{\mu_j, \phi_j, \sigma_j^2\}$ sequentially for each j as in Algorithm 1, we consider the following batch-wise update: update $\{\sigma_j^2\}_{j=1}^p$, then update $\{\mu_j\}_{j=1}^p$, and finally update $\{\phi_j\}_{j=1}^p$. Given $\{\mu_j, \phi_j, \sigma_j^2\}_{j=1}^p$, we can update $(\hat{\theta}, \hat{\sigma}^2)$ using Eq (6) and Eq (7).

2.4.1 Updating Equations

Update $\{\sigma_j^2\}_{j=1}^p$. We update σ_j^2's by maximizing

$$\begin{aligned}& \mathbb{E}^{q_1,\ldots,q_p}\left\{\sum_{i=1}^n\left[-\frac{(y_i - \mathbf{x}_i^T\boldsymbol{\beta})^2}{2\sigma^2}\right] + \sum_{j=1}^p \gamma_j\left[-\frac{\beta_j^2}{2v_1\sigma^2} - \log[f_j(\beta_j)]\right]\right\} \\ \propto & \sum_{j=1}^n -\phi_j\frac{n}{2\hat{\sigma}^2}\sigma_j^2 - \phi_j\frac{1}{2v_1\hat{\sigma}^2}\sigma_j^2 + \frac{1}{2}\phi_j\log(\sigma_j^2),\end{aligned}$$

and the updating equation for σ_j^2 is

$$\sigma_j^2 = \frac{\hat{\sigma}^2}{n + \frac{1}{v_1}}, \quad j = 1, \ldots, p. \tag{10}$$

Note that σ_j^2's take the same form for all j. The term n in the denominator is due to the fact that each column of $\mathbf{X}$ has been pre-processed such that $\|\mathbf{X}_j\|^2 = n$. Later in Section 3, in light of the asymptotic analysis, we will suggest to replaced n by $a_n \asymp n^a$ as in Eq (20).

Update $\{\mu_j\}_{j=1}^p$. We update μ_j's by maximizing

$$\mathbb{E}^{q_1,\ldots,q_p}\left\{\sum_{i=1}^n\left[-\frac{(y_i-\mathbf{x}_i^T\boldsymbol{\beta})^2}{2\sigma^2}\right]+\sum_{j=1}^p\gamma_j\left[-\frac{\beta_j^2}{2v_1\sigma^2}-\log[f_j(\beta_j)]\right]\right\},$$

and the updating equation for $\boldsymbol{\mu}$ is

$$\boldsymbol{\mu}=\left(\boldsymbol{\Phi}\mathbf{X}^T\mathbf{X}\boldsymbol{\Phi}+\boldsymbol{\Delta}+\frac{1}{v_1}\boldsymbol{\Phi}\right)^{-1}\boldsymbol{\Phi}\mathbf{X}^T\mathbf{y}, \tag{11}$$

where $\boldsymbol{\Delta}=\text{diag}\{\mathbf{X}^T\mathbf{X}\}\boldsymbol{\Phi}(\mathbf{I}-\boldsymbol{\Phi})=n\boldsymbol{\Phi}(\mathbf{I}-\boldsymbol{\Phi})$.

Update $\{\phi_j\}_{j=1}^p$. The objective function at this step involves a quadratic form of $\boldsymbol{\phi}=(\phi_1,\ldots,\phi_p)^T$. To simplify the computation, we apply a linear approximation to replace the quadratic term. The detailed derivation can be found in Appendix A, and the final updating equation for ϕ_j is given by

$$\begin{aligned}\text{Logit}(\phi_j)=&\,\text{Logit}(\hat{\theta})+\frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2}+\frac{\mu_j^2}{2\hat{\sigma}^2}\left(n+\frac{1}{v_1}\right)\\=&\,\text{Logit}(\hat{\theta})+\frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2}+\frac{\mu_j^2}{2\sigma_j^2}.\end{aligned} \tag{12}$$

Truncate ϕ_j's. The final expression of ϕ_j involves the exponential of $\text{Logit}(\phi_j)$, which could trigger the error of numerical overflow when the magnitude of the Logit value is large. In our implementation, we truncate the logit value, or equivalently restrict $\phi_j^{(t)}\in[c,1-c]$, where $1>c>0$ is a small constant.

We also stop updating any ϕ_j's once they reach the extreme values, c or $1-c$. That is, for $t>1$

$$\phi_j^{(t)}=\begin{cases}\text{Logit}^{-1}\left(\log\frac{\hat{\theta}}{1-\hat{\theta}}+\frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2}+\frac{\mu_j^2}{2\sigma_j^2}\right), & \text{if } \min\left(\phi_j^{(t-1)},1-\phi_j^{(t-1)}\right)>c;\\ \phi_j^{(t-1)}, & \text{if } \min\left(\phi_j^{(t-1)},1-\phi_j^{(t-1)}\right)\le c.\end{cases} \tag{13}$$

This stop-early updating scheme can dramatically reduce the computation cost for our algorithm, as explained in Section 2.4.2.

Stopping Criterion. After we loop over all the parameters mentioned above, we need to decide when to stop. A natural choice is to stop when the change of the objective function is less than some threshold. Since our primary focus is variable selection, we use the maximum entropy criterion: we compute the entropy for each $\text{Bern}(\phi_j)$, and stop if the maximum of the change of the entropy is less than a pre-specified threshold.

Algorithm 2 Batch-wise VB

Initialize $(\mu_1, \ldots, \mu_p)$, $(\sigma_1^2, \ldots, \sigma_p^2)$, $\hat{\theta}$, and $\hat{\sigma}^2$.

Initialize $(\phi_1, \ldots, \phi_p) = \mathbf{1}_{p\times 1}$, and truncation parameter c.

repeat

$\boldsymbol{\mu} \leftarrow (\boldsymbol{\Phi}\mathbf{X}^T\mathbf{X}\boldsymbol{\Phi} + \boldsymbol{\Delta} + \frac{1}{v_1}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}\mathbf{X}^T\mathbf{y}$

$\sigma_j^2 \leftarrow \frac{\hat{\sigma}^2}{a_n+1/v_1}$ or $\frac{\hat{\sigma}^2}{n+1/v_1}$

for j in $1:p$ **do**

if $\min(\phi_j, 1-\phi_j) > c$ **then**

$\phi_j \leftarrow \text{Logit}^{-1}\left\{\text{Logit}(\hat{\theta}) + \frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2} + \frac{\mu_j^2}{2\sigma_j^2}\right\}$

end if

end for

$\hat{\theta} \leftarrow \frac{\sum_{j=1}^p \phi_j + a_0 - 1}{p+a_0+b_0-2}$

$\hat{\sigma}^2 \leftarrow \frac{\|\mathbf{y}-\mathbf{X}\boldsymbol{\Phi}\boldsymbol{\mu}\|^2 + \sum_{j=1}^p[(\mathbf{X}_j^T\mathbf{X}_j(1-\phi_j)+1/v_1)\phi_j\mu_j^2 + (\mathbf{X}_j^T\mathbf{X}_j+1/v_1)\phi_j\sigma_j^2] + \frac{1}{v_1}\sum_{j=1}^p\phi_j(\mu_j^2+\sigma_j^2) + \nu\lambda}{n+\prod_{j=1}^p\phi_j+\nu+2}$

until Converge

return $(\mu_1, \ldots, \mu_p)$, $(\sigma_1^2, \ldots, \sigma_p^2)$, $(\phi_1, \ldots, \phi_p)$, $\hat{\theta}$, and $\hat{\sigma}^2$

2.4.2 Computational Complexity

The main computational cost lies in reverting a $p \times p$ matrix in Eq (11). The direct computation involves $O(p^3 \bigwedge n^3)$ operations which is time-consuming when both p and n are large. Next we describe the computation trick used in our implementation, which can dramatically reduce the computation cost.

At iteration t, Eq (11) can be rewritten as

$$\begin{aligned}\boldsymbol{\mu} &= \left(\mathbf{X}^T\mathbf{X}\boldsymbol{\Phi}^{(t)} + n(\mathbf{I}-\boldsymbol{\Phi}^{(t)}) + \frac{1}{v_1}\mathbf{I}\right)^{-1}\mathbf{X}^T\mathbf{y} \\ &= \left(A_{(t-1)} + \left(A_{(t)} - A_{(t-1)}\right)\right)^{-1}\mathbf{X}^T\mathbf{y},\end{aligned}$$

where $A_{(t)} = \left(\mathbf{X}^T\mathbf{X}\boldsymbol{\Phi}^{(t)} + n(\mathbf{I}-\boldsymbol{\Phi}^{(t)}) + \frac{1}{v_1}\mathbf{I}\right)$ and

$$A_{(t)} - A_{(t-1)} = \left(\mathbf{X}^T\mathbf{X} - n\mathbf{I}_p\right)\left(\boldsymbol{\Phi}^{(t)} - \boldsymbol{\Phi}^{(t-1)}\right).$$

At iteration $t > 1$, we would have $A_{(t-1)}^{-1}$ in hand. If the rank of $A_{(t)} - A_{(t-1)}$ is lower than p or n, then the problem can be reformulated as inverting a matrix of a lower rank.

Write $\mathbf{X}^T\mathbf{X} - n\mathbf{I}_p = B$ and $D^{(t)} = \boldsymbol{\Phi}^{(t)} - \boldsymbol{\Phi}^{(t-1)}$. Then $A_t - A_{(t-1)} = BD^{(t)}$. Based our experience, after the first several iterations, many ϕ_j's are close to 1 or 0, i.e., they will not be updated any more according to Eq (13). So many diagonal elements of $D^{(t)}$ are zero. Without loss of generality, assume only the first q elements of $D^{(t)}$ are not zero.

Then we have $BD^{(t)} = UCV$, where $U = B_{[,1:q]}$ contains only the first q columns of B, $C = \text{diag}\{\phi_j^{(t)} - \phi_j^{(t-1)}\}_{j=1}^q$, and $V = (\mathbf{I}_q, \mathbf{0}_{q\times(p-q)})$. Applying the Woodbury formula, we have

$$\boldsymbol{\mu} = \left(A_{(t-1)}^{-1} - A_{(t-1)}^{-1} U \left(C^{-1} + V A_{(t-1)}^{-1} U \right)_{q\times q}^{-1} V A_{(t-1)}^{-1} \right) \mathbf{X}^T \mathbf{y}.$$

So we only need to invert a $q \times q$ matrix where q is much smaller than p and n.

3. Asymptotic Analysis

Assume the response $\mathbf{y}$ is generated from the normal linear regression model (2) with $\boldsymbol{\beta}^* \in \mathbb{R}^p$ being the true regression coefficients. Let $\boldsymbol{\gamma}^*$ denote the true model index, i.e., $\gamma_j^* = 1$ if $\beta_j^* \neq 0$ and $\gamma_j^* = 0$ if $\beta_j^* = 0$. Also define the true set of relevant variables as

$$S^* = \{j : \beta_j^* \neq 0\} = \{j : \gamma_j^* = 1\}.$$

In our analysis, the dimension $p = p_n$ is allowed to diverge with n, and therefore $\boldsymbol{\beta}^* = \boldsymbol{\beta}_n^*$, $\boldsymbol{\gamma}^* = \boldsymbol{\gamma}_n^*$ and $S^* = S_n^*$ may also vary with n.

We will show that Algorithm 2 achieves both the frequentist consistency and the Bayesian consistency. Recall that Algorithm 2 returns an estimate of the relevant variable set via

$$\hat{S}_n = \{j : \phi_j > 0.5\}.$$

The frequentist consistency refers to the convergence (in probability) of $\hat{S}_n$ to S^*. The cut-off value 0.5 can be changed to any other fixed value in $(0, 1)$, which will not affect the consistency result as shown in our analysis.

In addition to a point estimate of the true variable set, our algorithm also returns a probability distribution over all 2^p variable sets (or sub-models), namely,

$$q(\boldsymbol{\gamma}) = \prod_{j=1}^{p} \left(\phi_j\right)^{\gamma_j} \left(1 - \phi_j\right)^{1-\gamma_j}.$$

The aforementioned frequentist consistency corresponds to $q(\boldsymbol{\gamma}^*)$ is the largest, i.e., the truth model receives the largest posterior probability, while the Bayesian consistency requires $q(\boldsymbol{\gamma}^*)$ converges to 1 in probability, which is stronger than the frequentist consistency.

In addition to the ordinary regularity conditions, for our algorithm to achieve consistency, we need to let v_1, the prior variance on the non-zero β_j's as in Eq (3), to grow to infinity at a certain rate of n. A similar condition also arises in the asymptotic study on Bayesian variable selection by Narisetty and He (2014) although their prior specification is different from ours. To help the readers to understand this condition, we first give the asymptotic analysis on a simple orthogonal design and then describe the general result.

3.1 The Orthogonal Design

Consider a simple case in which the design matrix is orthogonal, i.e., $\mathbf{X}^T\mathbf{X} = n$. To simplify our discussion, we also assume that σ^2 is known, θ is set to be $1/2$, and the minimal non-zero coefficient is bigger than some constant (i.e., the non-zero coefficients will not diminish to zero). These conditions will be relaxed in our result for the general case.

Suppose we run our algorithm for one step. From the updating equations of Algorithm 2, we have

$$2\,\mathrm{Logit}(\phi_j) = -\log(v_1 n + 1) + \frac{n\hat{\beta}_j^2}{\sigma^2}\frac{n}{n + \frac{1}{v_1}}, \tag{14}$$

where $\hat{\beta}_j = \frac{1}{n}\mathbf{X}_j^T\mathbf{y} \sim \mathrm{N}\left(\beta_j^*, \frac{\sigma^2}{n}\right)$ is the OLS estimator of the j-th coefficient.

When p is fixed, as in the classical asymptotic setting, it is easy to show that our algorithm has the desired asymptotic behavior as long as

$$v_1 n \to \infty, \quad \log(v_1 n) = o(n). \tag{15}$$

This is because: when $\beta_j^* = 0$, since $n\hat{\beta}_j^2 = O_P(1)$, the leading term in (14) is the first term that goes to $-\infty$, therefore ϕ_j goes to 0; when $\beta_j^* \neq 0$, the leading term in (14) is the second term that goes to ∞, therefore ϕ_j goes to 1.

When $p = p_n$ increases with n, the coefficients $(\beta_j^*)_{j=1}^p$ and the true variable set S_n^* may vary with n. As such, it is no longer meaningful to discuss the limit of Eq (14). Instead, we need to examine the limiting behavior of $\max_{j\notin S_n^*}\mathrm{Logit}(\phi_j)$ and $\min_{j\in S_n^*}\mathrm{Logit}(\phi_j)$.

First we show that the frequentist variable selection consistency could be achieved with $p = O(\sqrt{v_1 n})$, in addition to condition (15). Let C be an arbitrary positive number. It suffices to show that

$$P\left(\max_{j\notin S_n^*}\mathrm{Logit}(\phi_j) > -\frac{C}{2}\right) + P\left(\min_{j\in S_n^*}\mathrm{Logit}(\phi_j) < \frac{C}{2}\right) \to 0. \tag{16}$$

By the Bonferroni correction and the tail probability inequality of the normal distribution, we have

$$\begin{aligned} & P\left(\max_{j\notin S_n^*}\mathrm{Logit}(\phi_j) > -\frac{C}{2}\right) \\ \leq & \sum_{j\notin S_n^*} P\left(\frac{1}{\sigma^2}n\hat{\beta}_j^2\frac{n}{n+\frac{1}{v_1}} > \log(v_1 n + 1) - C\right) \\ \leq & \frac{c}{\sqrt{\log(v_1 n) - C}}\exp\left\{-\frac{\log(v_1 n) - C}{2} + \log p\right\} \to 0, \end{aligned} \tag{17}$$

as long as $p = O(\sqrt{v_1 n})$ and $v_1 n \to \infty$, where c is some constant. Similarly, we have

$$\begin{aligned}
&P\left(\min_{j\in S_n^*} \text{Logit}(\phi_j) < \frac{C}{2}\right) \le P\left(\frac{1}{2}\min_{j\in S_n^*} \frac{1}{\sigma^2} n\hat{\beta}_j^2 < \log(v_1 n + 1) + C\right) \\
\le& P\left(\frac{1}{\sigma}\sqrt{n}\min_{j\in S_n^*} |\beta_j^*| - \frac{1}{\sigma}\sqrt{n}\max_{j\in S_n^*} |\hat{\beta}_j - \beta_j^*| < \sqrt{2\log(v_1 n + 1) + 2C}\right) \\
\le& |S_n^*| P\left(\frac{1}{\sigma}\sqrt{n}|\hat{\beta}_j - \beta_j^*| > r\sqrt{n}\right),
\end{aligned}$$

which also goes to zero by the tail probability of the normal distribution, where r is some constant.

Secondly, we show that the Bayesian consistency could be achieved with $p = o(\sqrt{v_1 n})$. We can show that, if $p = o(\sqrt{v_1 n})$, Eq (16) still holds with a varying constant $C_n = s\log(v_1 n)$ where $s \in (0, 1)$. Then with probability going to 1, we have

$$\max_{j\in S_n^*}(1 - \phi_j) \bigvee \max_{j\notin S_n^*} \phi_j < \frac{1}{\exp\{\frac{1}{2}C_n\}} < \frac{1}{(v_1 n)^{s/2}}.$$

Using the inequality $\prod_j (1 - p_j) \ge 1 - \sum_j p_j$, we have

$$\begin{aligned}
q_(\boldsymbol{\gamma}^*) =& \Big[\prod_{j:\gamma_j^*=1} \phi_j\Big]\Big[\prod_{j:\gamma_j^*=0} (1 - \phi_j)\Big] \ge 1 - \sum_{j\in S_n^*}(1 - \phi_j) + \sum_{j\notin S_n^*} \phi_j \\
\ge& 1 - \frac{p}{(v_1 n)^{s/2}} \to 1.
\end{aligned}$$

Our analysis for the orthogonal case indicates that the choice of v_1 affects how fast we could allow p to diverge with n. For example, if v_1 is a constant, then our algorithm can achieve the frequentist consistency for $p = O(\sqrt{n})$ and Bayesian consistency for $p = o(\sqrt{n})$. In order to achieve consistency for larger p, we need to let $v_1 \to \infty$ with n.

3.2 The General Case

Without loss of generality, assume $S_n^* = \{1, ..., q_n\}$ and the true coefficient $\boldsymbol{\beta}^* = (\boldsymbol{\beta}_1^{*T}, 0^T)^T$, i.e, the first q_n features are the relevant ones. Write the design matrix as $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2)$ accordingly, where $\mathbf{X}_1$ is the $n \times q_n$ matrix corresponding to the signal features and $\mathbf{X}_2$ is the $n \times (p - q_n)$ matrix corresponding to the noise features.

In our analysis, we assume the following conditions hold.

(C1) Condition on model identification: Denote H_1 and H_2 as the projection matrices of $\mathbf{X}_1$ and $\mathbf{X}_2$ onto their column spaces respectively, then assume the rank of H_1 is q_n and the spectral norm of $H_1 H_2$ is upper bounded by 1, i.e., $\|H_1 H_2\|_2 < 1$.

(C2) Condition on the design matrix: Let λ_{n1} denote the minimal non-zero eigenvalue of matrix $\mathbf{X}^t\mathbf{X}$ and it satisfies

$$\lambda_{n1}^{-1} = O(n^{-\eta_1}), \quad 0 < \eta_1 \leq 1.$$

(C3) Condition on the sparsity of $\boldsymbol{\beta}^*$: The L_2 norm of the true regression coefficient $\boldsymbol{\beta}^*$ satisfies the following sparsity condition

$$\|\boldsymbol{\beta}^*\|_2^2 = O(n^{\eta_2}), \quad 0 \leq \eta_2 < \eta_1.$$

(C4) The beta-min condition: The minimal non-zero coefficient satisfies

$$\liminf_n \frac{\min_{j \in S_n^*} \sqrt{n}|\beta_j^*|}{n^{\eta_3/2}} \geq M, \tag{18}$$

where $\eta_3 \in (1 - \eta_1, 1]$ and $M > 0$ are two constants not depending on n.

(C5) Condition on the initial values: the initial value for all the inclusion probability $\phi_j^{(0)}$'s should be set to be 1, i.e., we start our algorithm with *all the variables in*. The initial values for the error variance $\hat{\sigma}^2$ and the Bernoulli parameter $\hat{\theta}$ could be any constants satisfying $0 < \hat{\sigma}^2 < \infty$ and $0 < \hat{\theta} < 1$. For the proof, we set

$$\phi_j^{(0)} = 1, \quad \hat{\sigma}^{2(0)} = 1, \quad \hat{\theta}^{(0)} = \frac{1}{2}. \tag{19}$$

In general, it is not realistic to derive consistent variable selection procedures and parameter estimation when the design matrix $\mathbf{X}$ is not of full rank (Shao and Deng, 2012). This is because the true coefficient $\boldsymbol{\beta}^*$ could be any vector from the set $\mathcal{B} = \{\boldsymbol{\beta} : \mathbf{X}\boldsymbol{\beta} = \mathbf{X}\boldsymbol{\beta}^*\}$ due to the collinearity among the columns of $\mathbf{X}$. Condition (C1) ensures that the true sparse coefficient $\boldsymbol{\beta}^*$ is identifiable. Let $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \boldsymbol{\beta}_2^T)^T$ be any vector from $\mathcal{B}$. Then we have $\mathbf{X}_1\boldsymbol{\beta}_1^* = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2$. Meanwhile (C1) implies that $\mathbf{X}_1\boldsymbol{\beta}_1^* = \mathbf{X}_1\boldsymbol{\beta}_1$. So $\|\boldsymbol{\beta}\|_2 \geq \|\boldsymbol{\beta}^*\|_2$ with equality holds true only if $\boldsymbol{\beta} = \boldsymbol{\beta}^*$. That is, the true coefficient vector $\boldsymbol{\beta}^*$ is the one from $\mathcal{B}$ with the smallest L_2 norm, which is unique.

In our algorithm, we approximate the posterior distribution of β_j by a mixture of a point mass at zero and a normal distribution. The updating equation (10) implies that the posterior variance σ_j^2 for non-zero β_j's is of order $1/n$, a reasonable result in the classical asymptotic setting where p is fixed and the minimal eigenvalue of $\mathbf{X}^T\mathbf{X}$ is of $O(n)$. However, in the diverging p case, as indicated by (C2), the minimal non-zero eigenvalue of $\mathbf{X}^T\mathbf{X}$ could be of order $O(n^{\eta_1})$. Therefore the posterior variance σ_j^2 defined in Eq (10) would be too small. In other words, we are too optimistic about the uncertainty of β_j. This is a common issue with variational algorithms, as pointed out in Bishop (2006) and shown in Figure 1:

the variational distribution tends to have a smaller support than the true target distribution. To fix this problem, we need to correct the posterior variance at the right order as follows:

$$\sigma_j^2 = \frac{\hat{\sigma}^2}{a_n + \frac{1}{v_1}}, \tag{20}$$

where $a_n \asymp n^a$ with $1 - \eta_3 < a < \eta_1$.

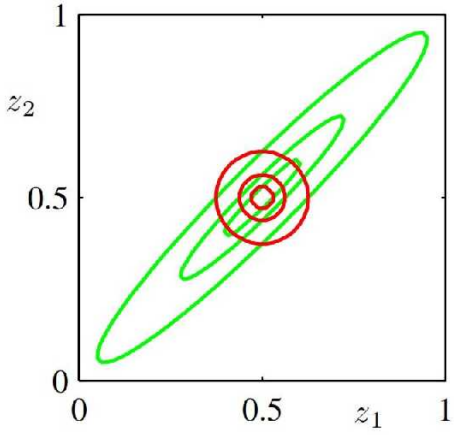

Figure 1: Page 468 from Bishop (2006): The green contours correspondens to the 1, 2, and 3 standard deviations for a correlated two-dimensional Gaussian distribution $p(z_1, z_2)$, and the red contours represent the corresponding levels for an approximating distribution $q_1(z_1)q_2(z_2)$ where q_1 and q_2 are obtained by minimization of the Kullback-Leibler divergence $KL(q\|p)$.

Below we summarize conditions on various rate parameters appearing in the our assumptions:

$$0 \leq \eta_2 < \eta_1 \leq 1, \tag{21}$$

$$0 \leq 1 - \eta_3 < a < \eta_1 \leq 1, \tag{22}$$

where (21) also appears in other work on variable selection, such as Shao and Deng (2012), and (22) indicates that the magnitude of the posterior variance (of order of $1/n^a$) should be between the minimal signal (of order $1/n^{1-\eta_3}$) and the maximal noise level (of order $1/n^{\eta_1}$). In the classical asymptotic setting, we have $\eta_1 = 1, \eta_2 = 0$ and $1 - \eta_3 = 0$.

With the modified σ_j^2 and some proper choice of v_1, we can show that our algorithm achieves the desired asymptotic property. We first present a lemma that shows that when the sample size is large enough, after one step of Algorithm 2, there is a gap between $\max_{j \notin S_n^*} \text{Logit}(\phi_j)$ and $\min_{j \in S_n^*} \text{Logit}(\phi_j)$ which is large enough for us to separate the relevant variables from the irrelevant ones. Using this lemma, we can then prove the frequentist

consistency and the Bayesian consistency. We present our asymptotic results below and include the proofs in Appendix B.

Lemma 1 *Assume conditions (C1-C5). Suppose v_1 is chosen such that*

$$\left(n^{-a} \bigvee n^{-(2\eta_1 - a - \eta_2)/2}\right) \prec v_1 \prec e^{n^{(a+\eta_3-1)}}$$

and $p = p_n$ satisfies $\log p_n = o(n^{\eta_1 - a})$, then for any constant $C > 0$, after one step of Algorithm 2, we have

$$P\left(\max_{j \notin S_n^*} \text{Logit}(\phi_j) > -\frac{C}{2}\right) \longrightarrow 0, \tag{23}$$

$$P\left(\min_{j \in S_n^*} \text{Logit}(\phi_j) < \frac{C}{2}\right) \longrightarrow 0. \tag{24}$$

Theorem 2 *(Variable Selection Consistency) Assume all conditions in Lemma 1, then we have*

$$P(\hat{S}_n = S_n^*) \longrightarrow 1.$$

Theorem 3 *(Bayesian Consistency) Assume conditions (C1-C5). Suppose v_1 is chosen such that*

$$\frac{p^2}{a_n} \prec v_1 \prec e^{n^{(a+\eta_3-1)}}$$

and $p = p_n$ satisfies $\log(p_n) = o(n^{\frac{a+\eta_3-1}{2}} \bigwedge n^{\eta_1 - a})$, then we have

$$q(\boldsymbol{\gamma}^*) = \left[\prod_{j:\gamma_j^*=1} \phi_j\right] \left[\prod_{j:\gamma_j^*=0} (1-\phi_j)\right] \xrightarrow{P} 1.$$

4. Simulation Studies

In this section, we conduct several simulation studies to compare the two VB algorithms: Algorithm 2 with Algorithm 1, and some other commonly used methods, like LASSO and SCAD.

In both VB algorithms, the choice of v_1 is chosen by cross-validation, and a sparse estimate of $\boldsymbol{\beta}$ is given by

$$\hat{\beta}_j = \mu_j \text{ if } \phi_j \geq 0.5; \quad 0 \text{ if } \phi_j < 0.5. \tag{25}$$

The value of a_n in the new variance update formula (20) is set to be λ_{n1}, the minimal non-zero eigenvalue of the sample covariance matrix. In our asymptotic analysis we set a_n to be of a smaller order of λ_{n1}; the main purpose of this choice is to ensure that the dimension p_n can grow exponentially fast, i.e., $\log p_n = o(n^{\eta_1 - a})$. In practice we have found that setting a_n to be λ_{n1} work well, along with the adaptive choice if v_1 via cross-validation.

4.1 Example 1: Benchmark Data

This is a popular benchmark data set, initially designed by Tibshirani (1994) and latter used by Fan and Li (2001) to compare different variable selection methods. The true coefficient is $\boldsymbol{\beta}^* = (3, 1.5, 0, 0, 2, 0, 0, 0)^T \in \mathbb{R}^8$, and the covariance between the i-th and the j-th variable is $0.5^{|i-j|}$. Denote the sample size by n, and the standard deviation of the error term by σ, we consider three scenarios: (1) $n = 40$ and $\sigma = 3$; (2) $n = 40$ and $\sigma = 1$; and (3) $n = 60$ and $\sigma = 1$. We repeat the simulation for 100 times and compare the results of Algorithm 1 and Algorithm 2 with the results of LASSO, SCAD, and the Oracle model from Fan and Li (2001).

To evaluation the estimation accuracy, we compute the model error (ME) by

$$\text{ME} = (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)^T \Sigma (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*)/\sigma^2 \tag{26}$$

where Σ is the covariance matrix of the eight covariates. We set the ME from the ordinary least square (OLS) model as the benchmark, and compute the relative model errors: dividing ME of the other models by that of the OLS. To obtain a robust criterion, we take the median of the relative model errors, namely the median of the relative model error (MRME). The results are shown in Table 1. When n gets larger or σ gets smaller, the two VB algorithms and SCAD become much better in terms of MRME, while Lasso does not gain obvious improvement. Overall, the two VB algorithms have lower MRME than SCAD. When the sample size gets larger and the noise level gets smaller, the MRMEs of the two VB algorithms are approaching to that of the Oracle. Also, Algorithm 2 is consistently better than Algorithm 1, especially when the sample size is small and the noise level is high.

To evaluation the selection accuracy, we count the number of zero coefficients among the signal and the noise variables. The results are also reported in Table 1. The "Correct"/"Incorrect" column records the average number of zero-coefficients returned by the method among noise/signal variables. A good method should have a number close to 5 for the "Correct" column and 0 for the "Incorrect" column. First, we notice that there is a trade-off between "Correct" and "Incorrect," or namely the trade-off between sensitivity and specificity of identifying noise variables. When the sample size is small ($n = 40$) and the noise level is high ($\sigma = 3$), Lasso identifies almost all the true predictors at the cost of including around 1.5 fake ones, while the variational Bayesian methods, on the other hand, correctly identify most noise variables at the cost of missing some signal variables. Hence, under such a low sample size and high variance setup, it is hard to tell which method is significantly better than the others. But when the sample size gets larger and/or the noise gets smaller, the two VB algorithms outperform all the other methods except the Oracle.

Method	MRME(%)	Avg. No. of 0 Coefficients	
		Correct	Incorrect
$n = 40, \sigma = 3$			
SCAD	72.90	4.20	0.21
Lasso	63.19	3.53	0.07
Alg.1	64.23	4.36	0.32
Alg.2	60.27	4.40	0.30
Oracle	33.31	5	0
$n = 40, \sigma = 1$			
SCAD	54.81	4.29	0
Lasso	63.19	3.51	0
Alg.1	39.34	4.85	0
Alg.2	37.37	4.92	0.12
Oracle	33.31	5	0
$n = 60, \sigma = 1$			
SCAD	47.54	4.37	0
Lasso	65.22	3.56	0
Alg.1	35.69	4.92	0
Alg.2	34.74	4.91	0
Oracle	29.82	5	0

Table 1: MRME and Average Number of 0 Coefficients. "Correct" represents how many noise variables are correctly identified, "Incorrect" means the number of true predictors being erroneously set to zero. The results for LASSO, SCAD, and Oracle are from Fan and Li (2001).

4.2 Example 2: Highly Correlated Noisy Data

This example is from Wang et al. (2011), in which the design matrix contains some highly correlated predictors with different signs. We use this example to demonstrate the advantage of the proposed batch-wise update used by Algorithm 2 over the component-wise update used by Algorithm 1: errors tend to accumulate with the component-wise update, especially when predictors are highly correlated.

The regression model has $p = 40$ highly correlated covariates and the true coefficient vector is $\boldsymbol{\beta}^0 = (3, 3, -2, 3, 3, -2, 0, \ldots, 0)^T$ with the last 34 elements being zero. The covariates are generated from a multivariate normal distribution: the variance of each variable is 1; the pairwise correlation of the first three variables is 0.9, that of the next three variables is 0.9, and all the other pairwise correlations are 0. The error terms are i.i.d. $N(0, 6^2)$. The sample size n is either 50 or 100. For each n, we repeat the experiment 100 times and compute ME that is defined in (26).

In Table 2, we report the average and the standard error of ME over 100 simulations. From the table, we can see that Algorithm 2 is better than Algorithm 1 for both $n = 50$ and $n = 100$. and it outperforms Lasso when the sample size grows.

	OLS	Lasso	Algorithm 1	Algorithm 2
$n = 50$	4913	233	322	305
	(323)	(11)	(17)	(18)
$n = 100$	706	144	139	109
	(25)	(6)	(7)	(7)

Table 2: $1000\times$ Average ME of Different Methods. The numbers in the parentheses are the corresponding $1000\times$ standard errors. The results for LASSO are from Wang et al. (2011).

In Table 3, we list the minimum, the median, and the maximum of the selection frequencies in pairs of parentheses of the 6 important variables (IV) and the 34 unimportant variables (UV). Although Algorithm 1 has the highest value for the maximum selection frequency of IV's, its median selection frequency of IV's is always the lowest. When the sample size increases, its median selection frequency of IV's even drops. Overall, Algorithm 2 is better than Algorithm 1.

In Table 4 and 5, we compare the estimated coefficients, as well as their signs, from Lasso, Algorithm 1, and Algorithm 2. For $n = 50$, the three methods can hardly detect any negative signs of β_3 or β_6: the high correlations between β_2 and β_3, and between β_5 and β_6 makes it difficult to identify the opposite signs of the neighboring highly-correlated

	Lasso	Algorithm 1	Algorithm 2
$n = 50$			
IV	(11, 70, 77)	(24, 41 ,82)	(24,61,71)
UV	(12, 17, 25)	(5, 11,17)	(7,11,19)
$n = 100$			
IV	(8, 84, 88)	(12, 33, 99)	(18, 68, 73)
UV	(12, 22, 31)	(4, 7, 10)	(0, 0, 2)

Table 3: Variable Selection Frequencies (%). IV: important variables; UV: unimportant variables. The three numbers in parentheses are the min, median, and max of selection frequencies among all important or unimportant variables, respectively. The results for LASSO are from Wang et al. (2011).

predictors. When $n = 100$, Algorithm 2 performs the best. Firstly, it can identify the negative signs of β_3 and β_6 in some simulations, while Lasso and Algorithm 1 can barely identify any negative signs of β_3 and β_6. Secondly, the magnitude of the estimates from Algorithm 2 is closer to the true coefficients, and the average coefficient estimates of β_3 and β_6 are negative, which is correct. Estimates from Lasso seem to deviate the most from the truth, which can also be confirmed by Table 2.

4.3 Example 3: Large p, Small n

We first consider a large-p-small-n example from Ročková and George (2014), in which only the first three of $p = 1000$ predictors are the true ones with non-zero coefficients to be 3, 2, and 1 respectively, and the sample size $n = 100$. The covariance between the i-th and j-th variables is $0.6^{|i-j|}$, and the error terms are generated from $i.i.d.$ $N(0, 3)$.

We fit the model using Algorithm 1 and Algorithm 2, with fixed $v_1 = 1$. As shown in Figure 2a and 2b, both algorithms make no mistake in variable selection and the estimates are very close to the true coefficients.

Next we make this example a little more challenging by adding more non-zero coefficients of different magnitudes. The true coefficient vector is now set to be $\boldsymbol{\beta}^0 = (\boldsymbol{\beta}_1^{0T}, 0, \ldots, 0)^T$, where the last 980 elements are all zero and the 20 non-zero coefficients in $\boldsymbol{\beta}_1^0$ contains randomly distributed ten 1's, seven 2's, and three 3's. With the 0.6 pairwise correlation, some weak signals may be overshadowed by nearby strong signals, which makes variable selection a challenging task. We repeat the experiment 100 times, and compare results from Algorithm 1, Algorithm 2 and Lasso. For the tuning parameter in Lasso, we use the

	β_1	β_2	β_3	β_4	β_5	β_6
True Coef	3	3	-2	3	3	-2
Lasso						
Ave. of est.	1.41	1.27	0.12	1.36	1.36	0.09
	(0.12)	(0.14)	(0.04)	(0.12)	(0.13)	(0.06)
No. of pos. sgn.	73	69	13	74	68	13
No. of neg. sgn.	0	0	1	0	0	2
Alg.1						
Ave. of est.	2.08	0.75	0.23	2.19	0.61	0.23
	(0.16)	(0.11)	(0.05)	(0.16)	(0.09)	(0.04)
No. of pos. sgn.	80	43	24	82	40	26
No. of neg. sgn.	0	0	0	0	0	0
Alg.2						
Ave. of est.	1.63	1.17	0.30	1.36	1.44	0.28
	(0.16)	(0.13)	(0.09)	(0.14)	(0.14)	(0.07)
No. of pos. sgn.	71	57	31	65	66	24
No. of neg. sgn.	0	0	1	0	0	0

Table 4: Coefficient and Coefficient Sign Estimation, $n = 50$. The numbers in parenthesis are the standard errors. The results for LASSO are from Wang et al. (2011).

	β_1	β_2	β_3	β_4	β_5	β_6
True Coef	3	3	-2	3	3	-2
Lasso						
Ave. of est.	1.67	1.50	0.06	1.85	1.38	0.04
	(0.11)	(0.11)	(0.02)	(0.13)	(0.13)	(0.03)
No. of pos. sgn.	91	85	8	86	78	9
No. of neg. sgn.	0	0	0	0	0	2
Alg.1						
Ave. of est.	2.97	0.56	0.08	2.93	0.56	0.10
	(0.12)	(0.08)	(0.02)	(0.12)	(0.09)	(0.03)
No. of pos. sgn.	99	34	12	96	32	12
No. of neg. sgn.	0	0	0	0	0	0
Alg.2						
Ave. of est.	2.09	1.92	-0.11	2.33	1.85	-0.20
	(0.16)	(0.16)	(0.05)	(0.18)	(0.18)	(0.09)
No. of pos. sgn.	73	66	5	72	65	7
No. of neg. sgn.	0	1	13	0	3	16

Table 5: Coefficient and Coefficient Sign Estimation, $n = 100$. The numbers in parenthesis are the standard errors. The results for LASSO are from Wang et al. (2011).

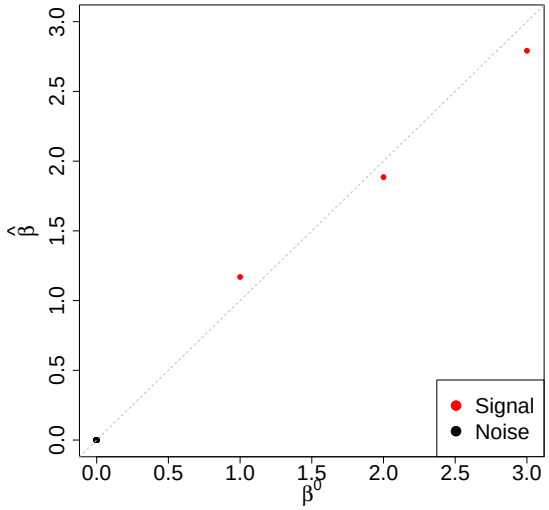

(a) True Coefficient vs. Sparse Estimates: Algorithm 1

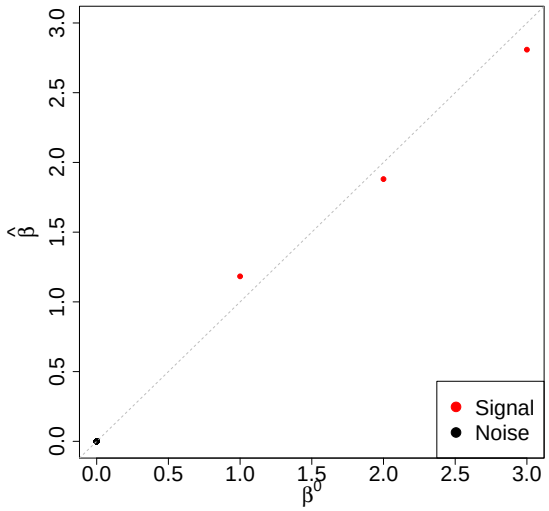

(b) True Coefficient vs. Sparse Estimates: Algorithm 2

Figure 2: $\hat{\boldsymbol{\beta}}$ *vs.* $\boldsymbol{\beta}^0$ for the large-p-small-n example from Ročková and George (2014).

default `cv.glmnet` function from the `R` package `glmnet` (Friedman et al., 2010), and report the results for both `lambda.min` that is the λ with the smallest CV error, and `lambda.1se` that is the largest λ whose CV error is within one standard error of the smallest CV error.

The results are summarized in Table 6. Regarding the accuracy for identifying important variables, Algorithm 1 is the worst: on average it only identifies 8.97 true predictors, fewer than a half. This is supported with our large sample analysis: with high correlation, the errors tend to accumulate, which makes some true predictors hard to be identified. Regarding the accuracy for identifying unimportant variables, `Lasso.min` is the worst: on average it selects 34.14 noise variables, more than doubled that of `Lasso.1se`, the second worst. Overall, Algorithm 2 performs the best.

	Model	Algorithm 1	Algorithm 2	Lasso.min	Lasso.1se
# of IV	Mean	8.97	15.61	19.92	19.77
	S.E.	(0.13)	(0.19)	(0.03)	(0.05)
# of UV	Mean	0.12	0.5	34.14	11.23
	S.E.	(0.04)	(0.08)	(1.50)	(0.81)

Table 6: Variable Selection Frequencies (%). IV: important variables; UV: unimportant variables. The numbers in parenthesis are the standard errors.

5. Real Data: Boston Housing Data

5.1 Introduction

The original data is from the `R` library `mlbench`, which has 506 observations on 19 variables. We apply some suggested transformations on the data according to Johnson et al. (1992), then remove three variables `medv`, `town`, and `tract`, and use `cmedv` as the response variable. We call this data set "Boston Housing 1."

Then we create a larger data set called "Boston Housing 2." First we add all the 119 quadratic terms (including all pairwise interaction terms) of the predictors, and 500 noise features. We generate the noise features in 50 batches. For each batch, we randomly select 10 variables from the set of 134 "true" variables, which gives us a 506×10 data matrix; for each entry of the matrix, we add a small Gaussian error, and then randomly shuffle the 506 rows. So each noise feature looks like some true variable marginally, and in addition correlations among the true variables are preserved in the noise features. The final data set has 634 features, which is larger than the sample size.

As we have shown that Algorithm 2 is better than Algorithm 1 in the simulation study, we just compare Algorithm 2 to Lasso, Ridge, and the full model (i.e., the model using all predictors). For Algorithm 2, we consider the following two prediction methods:

(i) Use the sparse estimate $\hat{\boldsymbol{\beta}}$ that is defined at (25) and the prediction is given by $\hat{\mathbf{y}} = \mathbf{X}^{new}\hat{\boldsymbol{\beta}}$. We abbreviate this approach by "S" that stands for "Sparsity Prediction."

(ii) Refit a linear regression model using the predictors from $\hat{S} = \{j : \phi_j > 0.5\}$ and denote the estimated coefficients from the OLS by $\hat{\boldsymbol{\beta}}^{ols}$. Then the prediction is given by $\hat{\mathbf{y}} = \mathbf{X}^{new}_{[,\hat{S}]}\hat{\boldsymbol{\beta}}^{ols}$, where $\mathbf{X}^{new}_{[,\hat{S}]}$ denotes a subset of the data matrix $\mathbf{X}^{new}$ with only columns from $\hat{S}$. We abbreviate this approach by "TS" that stands for "Two-Stage Sparse Prediction".

We run the following simulation for 50 times. In each iteration, we randomly subset 75% of the dataset as the training data (380 observations) and predict on the remaining validation data (126 observations). Then we record the selected model size and compute the mean squared prediction error (MSPE). As all the methods but the OLS have a tuning parameter, we select the tuning parameter using cross-validation. For Algorithm 2, we use a 5-fold cross validation for the two prediction approachs; for Lasso and ridge regression, we use the default `cv.glmnet` function from the `R` package `glmnet` (Friedman et al., 2010), and report results for both `lambda.min` and `lambda.1se`. The results are summarized in Table 7.

5.2 Boston Housing 1

Most methods perform similarly based on the prediction error. Surprisingly the full model performs the best. This is because the potential gain of variable selection for such a traditional small-p-large-n example is negligible. On the other hand, the potential bias introduced by variable selection or shrinkage procedures, when relevant variables are mistakenly excluded or over-shrunk, can be large. This also explains why Lasso.1se/ridge.1se performs worse than Lasso.min/ridge.min, since the former tends to pick a smaller model than the latter, i.e., has a higher chance of missing relevant variables.

5.3 Boston Housing 2

The ridge regression is the worst of all methods: both ridge.1se and ridge.min have relatively large effective dimensions, but high prediction errors. For prediction accuracy, Lasso.min, Alg2.S and Alg2.TS are better than the other methods; regarding sparsity, the models selected by Lasso.1se, Alg2.S and Alg2.TS are much smaller than the others. The best model is Alg2.TS: it has the best prediction accuracy with the most sparse model.

	Model	Full	Ridge.min	Ridge.1se	Lasso.min	Lasso.1se	Alg2.S	Alg2.TS
BH 1	Mean	0.043	0.044	0.049	0.044	0.048	0.045	0.044
	SE	(0.008)	(0.009)	(0.010)	(0.008)	(0.008)	(0.009)	(0.008)
	Size	15	12.34	9.21	13.84	7.2	10.8	11.36
BH 2	Mean		0.065	0.071	0.043	0.047	0.046	0.042
	SE		(0.011)	(0.012)	(0.008)	(0.008)	(0.012)	(0.007)
	Size		52.21	40.21	38.7	8.68	9.18	7.74

Table 7: Boston Housing Data: Average and S.E. of MSPE, and Model Size

6. Conclusions

The Bayesian approach to variable selection is appealing since it outputs not only a single model but a probability distribution over all possible models. Hence model uncertainty can be naturally incorporated into estimation, prediction, and many other statistical inferences. However, most Bayesian variable selection methods are implemented through MCMC, which is time consuming when the model dimension is large. In this paper, we propose an algorithm that approximates the posterior distribution via a variational optimization. Our proposed algorithm converges very fast and can scale up with large data sets. We also showed that the approximation returned by our algorithm has the desired asymptotic behavior, which achieves both the frequentist consistency and Bayesian consistency asymptotically.

Appendices

Appendix A. Update ϕ_j's in Algorithm 2

We provide the detailed derivation of the updating equation for ϕ_j's here. The derivation of other variational distributions and MAP estimators is straightforward.

We fix $\{\mu_j\}_{j=1}^p$, $\{\sigma_j^2\}_{j=1}^p$, $\hat{\theta}$, and $\hat{\sigma}^2$, and update $\{\phi_j\}_{j=1}^p$. The objective function is given by

$$\begin{aligned}\Omega(\phi_1,\ldots,\phi_p) = &-\frac{1}{2\hat{\sigma}^2}\left[\mathbf{y}^T\mathbf{y} - 2\mathbf{y}^T\mathbf{X}\mathbb{E}(\boldsymbol{\beta}) + \mathbb{E}\left(\boldsymbol{\beta}^T(\mathbf{X}^T\mathbf{X})\boldsymbol{\beta}\right)\right]\\ &+\sum_{j=1}^p \phi_j\left(\frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2} + \frac{1}{2} - \frac{\mu_j^2+\sigma_j^2}{2v_1\hat{\sigma}^2} + \log(\hat{\theta}) - \log(\phi_j)\right)\\ &+(1-\phi_j)[\log(1-\hat{\theta}) - \log(1-\phi_j)] + \text{Constant}.\end{aligned}$$

Denote $\mathbf{U} = \text{diag}\{\mu_1,\ldots,\mu_p\}$ and $\boldsymbol{\phi} = (\phi_1,\ldots,\phi_p)^T$, we have $\mathbb{E}(\boldsymbol{\beta}) = \mathbf{U}\boldsymbol{\phi}$ and

$$\begin{aligned}&\mathbb{E}\left(\boldsymbol{\beta}^T(\mathbf{X}^T\mathbf{X})\boldsymbol{\beta}\right)\\ =&\,\text{tr}\left((\mathbf{X}^T\mathbf{X})\,\text{Cov}(\boldsymbol{\beta})\right) + \mathbb{E}(\boldsymbol{\beta})^T(\mathbf{X}^T\mathbf{X})\mathbb{E}(\boldsymbol{\beta})\\ =&\sum_{j=1}^p n(\mu_j^2+\sigma_j^2)\phi_j - n\boldsymbol{\phi}^T\mathbf{U}\mathbf{U}\boldsymbol{\phi} + \boldsymbol{\phi}^T\mathbf{U}(\mathbf{X}^T\mathbf{X})\mathbf{U}\boldsymbol{\phi}.\end{aligned}$$

Then the objective function becomes

$$\begin{aligned}\Omega(\phi_1,\ldots,\phi_p) =&\frac{1}{\hat{\sigma}^2}\mathbf{y}^T\mathbf{X}\mathbf{U}\boldsymbol{\phi} - \frac{n}{2\hat{\sigma}^2}\sum_{j=1}^p(\mu_j^2+\sigma_j^2)\phi_j\\ &-\frac{1}{2\hat{\sigma}^2}\boldsymbol{\phi}^T\mathbf{U}(\mathbf{X}^T\mathbf{X} - n\mathbf{I})\mathbf{U}\boldsymbol{\phi}\\ &+\sum_{j=1}^p \phi_j\left(\frac{1}{2}\log\frac{\sigma_j^2}{v_1\hat{\sigma}^2} + \frac{1}{2} - \frac{\mu_j^2+\sigma_j^2}{2v_1\hat{\sigma}^2} + \log(\hat{\theta}) - \log(\phi_j)\right)\\ &+(1-\phi_j)[\log(1-\hat{\theta}) - \log(1-\phi_j)] + \text{Constant}.\end{aligned}$$

Taking derivative w.r.t. $\boldsymbol{\phi}$, we have

$$\begin{aligned}\nabla\Omega =&\frac{1}{\hat{\sigma}^2}\mathbf{y}^T\mathbf{X}\mathbf{U} - \frac{n}{2\hat{\sigma}^2}\{(\mu_j^2+\sigma_j^2)\}_{j=1}^p\\ &-\frac{1}{\hat{\sigma}^2}\mathbf{U}(\mathbf{X}^T\mathbf{X} - n\mathbf{I})\mathbf{U}\boldsymbol{\phi} + Rest.\end{aligned}$$

Direct optimization would involve numerical methods due to the non-linear system with constraints. Thus, we take an approximation approach. Denote $\Delta = \mathbf{U}(\mathbf{X}^T\mathbf{X} - n\mathbf{I}_p\})\mathbf{U}$

and define $g(\boldsymbol{\phi}) = \boldsymbol{\phi}^T \Delta \boldsymbol{\phi}$. At the t-th iteration, using the Taylor expansion, we approximate the quadratic form $g(\boldsymbol{\phi}^{(t)})$ by

$$\begin{aligned} g(\boldsymbol{\phi}^{(t)}) \approx & g(\boldsymbol{\phi}^{(t-1)}) + \nabla(g(\boldsymbol{\phi}^{(t-1)}))^T(\boldsymbol{\phi}^{(t)} - \boldsymbol{\phi}^{(t-1)}) \\ = & (\boldsymbol{\phi}^{(t-1)})^T \Delta \boldsymbol{\phi}^{(t-1)} + 2(\boldsymbol{\phi}^{(t-1)})^T \Delta(\boldsymbol{\phi}^{(t)} - \boldsymbol{\phi}^{(t-1)}) \\ \propto & 2(\boldsymbol{\phi}^{(t-1)})^T \Delta \boldsymbol{\phi}^{(t)}. \end{aligned}$$

Hence, we have

$$\begin{aligned} \Omega \approx & \frac{1}{\hat{\sigma}^2} \mathbf{y}^T \mathbf{X} \mathbf{U} \boldsymbol{\phi} - \frac{n}{2\hat{\sigma}^2} \sum_{j=1}^{p} (\mu_j^2 + \sigma_j^2) \phi_j \\ & - \frac{1}{\hat{\sigma}^2} (\boldsymbol{\phi}^{(t-1)})^T \mathbf{U} (\mathbf{X}^T \mathbf{X} - n\mathbf{I}) \mathbf{U} \boldsymbol{\phi} \\ & + \sum_{j=1}^{p} \phi_j \left(\frac{1}{2} \log \frac{\sigma_j^2}{v_1 \hat{\sigma}^2} + \frac{1}{2} - \frac{\mu_j^2 + \sigma_j^2}{2 v_1 \hat{\sigma}^2} + \log(\hat{\theta}) - \log(\phi_j) \right) \\ & + (1 - \phi_j)[\log(1 - \hat{\theta}) - \log(1 - \phi_j)] + \text{Constant}. \end{aligned}$$

Taking partial derivative w.r.t. ϕ_j's respectively, we have

$$\begin{aligned} \frac{\partial \Omega}{\partial \phi_j} \approx & \frac{1}{\hat{\sigma}^2} \mathbf{y}^T \mathbf{X}_j \mu_j - \frac{n}{2\hat{\sigma}^2} (\mu_j^2 + \sigma_j^2) - \frac{1}{\hat{\sigma}^2} \sum_{k \neq j} \mu_j \mu_k \mathbf{X}_k^T \mathbf{X}_j \phi_k^{(t-1)} \\ & + \frac{1}{2} \log \frac{\sigma_j^2}{v_1 \hat{\sigma}^2} + \frac{1}{2} - \frac{\mu_j^2 + \sigma_j^2}{2 v_1 \hat{\sigma}^2} + \log(\hat{\theta}) - \log(\phi_j) - 1 - \log(1 - \hat{\theta}) + \log(1 - \phi_j) + 1. \end{aligned}$$

Setting $\frac{\partial \Omega}{\partial \phi_j} = 0$, we have

$$\text{Logit}(\phi_j) = \text{Logit}(\hat{\theta}) + \frac{1}{2} \log \frac{\sigma_j^2}{v_1 \hat{\sigma}^2} - \frac{\mu_j^2}{2\sigma_j^2} + \frac{1}{\hat{\sigma}^2} \mu_j \mathbf{X}_j^T \left(\mathbf{y} - \mathbf{X}_{[,-j]} \boldsymbol{\Phi}_{[-j,-j]}^{(t-1)} \boldsymbol{\mu}_{[-j]} \right), \qquad (27)$$

where the last term is equal to $\frac{1}{\hat{\sigma}^2} \mu_j \mathbf{X}_j^T \mu_j (n + 1/v_1)$ according to Algorithm 1. Therefore, we further approximate it by $\frac{1}{\hat{\sigma}^2} \mu_j^2 (n + 1/v_1)$ to reduce computational complexity. Then we have

$$\begin{aligned} \text{Logit}(\phi_j) = & \text{Logit}(\hat{\theta}) + \frac{1}{2} \log \frac{\sigma_j^2}{v_1 \hat{\sigma}^2} + \frac{\mu_j^2}{2\hat{\sigma}^2} \left(n + \frac{1}{v_1} \right) \\ = & \text{Logit}(\hat{\theta}) + \frac{1}{2} \log \frac{\sigma_j^2}{v_1 \hat{\sigma}^2} + \frac{\mu_j^2}{2\sigma_j^2}. \end{aligned}$$

Appendix B. Proofs for Section 3

B.1 Proof for Lemma 1

In Algorithm 2, we first update $\{\mu_{nj}\}_{j=1}^p$, given the initial value $\phi_j = 1$. The updating formula for $\{\mu_{nj}\}_{j=1}^p$ at the first iteration is

$$\begin{aligned}\boldsymbol{\mu}_n &= \left(\boldsymbol{\Phi}\mathbf{X}^T\mathbf{X}\boldsymbol{\Phi} + n\boldsymbol{\Phi}(\mathbf{I}_p - \boldsymbol{\Phi}) + \frac{1}{v_1}\boldsymbol{\Phi}\right)^{-1}\boldsymbol{\Phi}\mathbf{X}^T\mathbf{y}_n \\ &= \left[\mathbf{X}^T\mathbf{X} + \frac{1}{v_1}\mathbf{I}_p\right]^{-1}\mathbf{X}^T(\mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\epsilon}_n),\end{aligned}$$

where $\boldsymbol{\Phi} = \mathbf{I}_p$. After updating $\boldsymbol{\mu}_n$, given the initial values $\hat{\theta} = 1/2$ and $\hat{\sigma}^2 = 1$, we update the logit of ϕ_j using

$$2\,\mathrm{Logit}(\phi_j) = -\log(a_n v_1 + 1) + \mu_{nj}^2\left(a_n + \frac{1}{v_1}\right).$$

To quantify the magnitude of $\mathrm{Logit}(\phi_j)$, the key is to quantify μ_{nj}. Decompose $\boldsymbol{\mu}_n$ into three parts: the true coefficient vector $\boldsymbol{\beta}^*$, the bias $\mathbf{b}_n$, and the error projection $\mathbf{w}_n$.

$$\begin{aligned}\boldsymbol{\mu}_n =& \boldsymbol{\beta}^* - \left(v_1\mathbf{X}^T\mathbf{X} + \mathbf{I}_p\right)^{-1}\boldsymbol{\beta}^* + \left(\mathbf{X}^T\mathbf{X} + \frac{1}{v_1}\mathbf{I}_p\right)^{-1}\mathbf{X}^T\boldsymbol{\epsilon}_n \\ =& \boldsymbol{\beta}^* - \mathbf{b}_n + \mathbf{w}_n.\end{aligned}$$

Next we prove the following results.

1. **Bound for $\mathbf{b}_n$.** Denote the singular value decomposition of $\mathbf{X}$ as PDQ^T, where r is the rank of $\mathbf{X}$, the dimension of P, D, and Q are $n \times r$, $r \times r$, and $p \times r$, respectively. Condition (C1) implies that the true coefficient vector $\boldsymbol{\beta}^*$ is the projection of the set $\mathcal{B}$ onto the row space of $\mathbf{X}$, i.e., $\boldsymbol{\beta}^* = QQ^T\boldsymbol{\beta}^*$. Then the bias term can be written as

$$\begin{aligned}\mathbf{b}_n &= \left(v_1\mathbf{X}^T\mathbf{X} + \mathbf{I}_p\right)^{-1}\boldsymbol{\beta}^* \\ &= (v_1 QD^2Q^T + \mathbf{I}_p)^{-1}QQ^T\boldsymbol{\beta}^* \\ &= Q(v_1D^2 + \mathbf{I}_r)^{-1}Q^T\boldsymbol{\beta}^*.\end{aligned}$$

So we can bound the maximal of the bias term b_{nj} by

$$\begin{aligned}\max_j b_{nj}^2 &\le \|\mathbf{b}_n\|_2^2 = \|Q(v_1D^2 + \mathbf{I}_r)^{-1}Q^T\boldsymbol{\beta}^*\|_2^2 \\ &\le \left(\frac{1}{v_1\lambda_{n1} + 1}\right)^2\|\boldsymbol{\beta}^*\|_2^2 = \frac{O(n^{\eta_2})}{(v_1\lambda_{n1})^2}. \qquad (28)\end{aligned}$$

2. **Bound for the variance of w_{nj}.** $\mathbf{w}_n$ is a Gaussian random variable with mean zero and covariance

$$\begin{aligned}\mathrm{Cov}(\mathbf{w}_n) =& \mathrm{Cov}((\mathbf{X}^T\mathbf{X} + \frac{1}{v_1}\mathbf{I}_p)^{-1}\mathbf{X}^T\boldsymbol{\epsilon}_n) \\ =& \sigma^2 Q \,\mathrm{diag}\left(\frac{d_j^2}{(d_j^2 + 1/v_1)^2}\right) Q^T\end{aligned}$$

where d_j's are the elements from the diagonal matrix D. So the variance of each w_{nj} is bounded by the largest eigenvalue of $\mathrm{Cov}(\mathbf{w}_n)$:

$$\mathrm{Var}(w_{nj}) \leq \frac{\sigma^2\lambda_{n1}}{(\lambda_{n1} + 1/v_1)^2} \leq \frac{\sigma^2}{\lambda_{n1}}. \tag{29}$$

3. **Inequalities for μ_{nj}^2.** Using the inequality $(\sum_{i=1}^n a_i)^2 \leq n\sum_{i=1}^n a_i^2$, we have

$$\begin{aligned}\max_{j\notin S_n^*} \mu_{nj}^2 &\leq 2(\max_j b_{nj}^2 + \max_j w_{nj}^2), \\ \min_{j\in S_n^*} \mu_{nj}^2 &\geq \frac{1}{3}\min_{j\in S_n^*}(\beta_j^*)^2 - \max_j b_{nj}^2 - \max_j w_{nj}^2.\end{aligned}$$

First we show (23). Note

$$\max_{j\notin S_n^*} 2\,\mathrm{Logit}(\phi_j) \leq -\log(a_n v_1 + 1) + 2\Big(a_n + \frac{1}{v_1}\Big)\Big(\max_j b_{nj}^2 + \max_j w_{nj}^2\Big).$$

We have

(a) $\log(v_1 a_n + 1) \to \infty$ since $v_1 \succ n^{-a}$.

(b) By (28), $(a_n + \frac{1}{v_1})\max_j b_{nj}^2 = O(n^{\eta_2+a-2\eta_1}/v_1^2) \to 0$ since $v_1 \succ n^{-\frac{1}{2}(2\eta_1-a-\eta_2)}$.

(c) For $(a_n + 1/v_1)w_{nj}^2$, its variance is upper bounded by $O(n^{a-\eta_1})$. Hence, for any constant c,

$$\begin{aligned}& P\left((a_n + 1/v_1)\max_j w_{nj}^2 > c\right) \\ \leq& \sum_{j=1}^p P\left((a_n + 1/v_1)w_{nj}^2 > c\right) \\ \leq& \frac{p}{\sqrt{2\pi n^{\eta_1-a}}}\exp\left\{-\frac{c}{2}n^{\eta_1-a}\right\} \longrightarrow 0,\end{aligned}$$

since $\log p = o(n^{\eta_1-a})$.

Thus, $\max_{j\notin S_n^*} 2\,\mathrm{Logit}(\phi_j) \xrightarrow{P} -\infty$ as $n\to\infty$, and therefore (23) holds true.

Next we show (24). Note

$$\min_{j\in S_n^*} 2\,\mathrm{Logit}(\phi_{nj}) \geq -\log(a_n v_1+1) + \Big(a_n+\frac{1}{v_1}\Big)\Big(\frac{1}{3}\min_{j\in S_n^*}(\beta_j^*)^2 - \max_j b_{nj}^2 - \max_j w_{nj}^2\Big).$$

Since $\Big(a_n+\frac{1}{v_1}\Big)\frac{1}{3}\min_{j\in S_n^*}(\beta_j^*)^2 \succeq n^{a-(1-\eta_3)} \succ \log(a_n v_1+1)$, it is the leading term. So $\min_{j\in S_n^*} 2\,\mathrm{Logit}(\phi_{nj}) \xrightarrow{P} \infty$ as $n\to\infty$ and therefore (24) holds true.

B.2 Proof for Theorem 2

With Lemma 1, it is easy to show that when the sample size is large enough, our algorithm will stop with one update. For any threshold value c in (13), we can set C in Lemma 1 to be bigger than $2\log(\frac{1}{c}-1)$. Then with probability going to 1, after the first iteration, we will have $\max_{j\notin S_n^*}\phi_j < c$, and $\min_{j\in S_n^*}\phi_{nj} > 1-c$ and the algorithm will halt. Therefore, $\hat{S}_n = S_n^*$ with probability 1 when $n\to\infty$. Hence, the frequentist selection consistency follows.

B.3 Proof for Theorem 3

Let $C_n = s\log(v_1 a_n)$, where $s\in(0,1)$. Following the same argument used in the proof of Lemma 1, we can show that

$$P\left(2\max_{j\notin S_n^*}\mathrm{Logit}(\phi_j) > -C_n\right)\to 0 \quad\text{and}\quad P\left(2\min_{j\in S_n^*}\mathrm{Logit}(\phi_j) < C_n\right)\to 0,$$

if $\log p = o(n^{\eta_1-a})$. Therefore, with probability going to 1, we have

$$\max_{j\notin S_n^*}\log\frac{\phi_j}{1-\phi_j} < -\frac{1}{2}C_n \Rightarrow \max_{j\notin S_n^*}\phi_j < \frac{1}{\exp\{\frac{1}{2}C_n\}},$$
$$\min_{j\in S_n^*}\log\frac{\phi_j}{1-\phi_j} > \frac{1}{2}C_n \Rightarrow \max_{j\in S_n^*}(1-\phi_j) < \frac{1}{\exp\{\frac{1}{2}C_n\}}.$$

Using the inequality $\prod_j(1-p_j)\geq 1-\sum_j p_j$, we have

$$1-q(\boldsymbol{\gamma}^*) \leq \sum_{j\in S_n^*}(1-\phi_j) + \sum_{j\notin S_n^*}\phi_j \leq p\times\left[\max_{j\in S_n^*}(1-\phi_j)\bigvee\max_{j\notin S_n^*}\phi_j\right] = \frac{p}{(v_1 a_n)^{s/2}}.$$

If $p\prec(v_1 a_n)^{s/2}$, we have

$$1-q(\boldsymbol{\gamma}^*) \leq \frac{p}{e^{\frac{1}{2}C_n}} = \frac{p}{(v_1 a_n)^{s/2}} \xrightarrow{P} 0, \tag{30}$$

as $v_1 \prec \exp(n^{a-(1-\eta_3)})$, $p\prec(v_1 a_n)^{s/2}\prec\exp(\sqrt{n}^{\frac{a-1+\eta_3}{2}})n^{s/2}$. This implies that if $p = o(\exp(n^{\frac{a-1-\eta_3}{2}}))$, we can achieve Bayesian consistency by letting v_1 going to infinity at an exponential order.

References

Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

Peter Bühlmann and Sara van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Publishing Company, Incorporated, 1st edition, 2011.

Peter Carbonetto and Matthew Stephens. Scalable variational inference for bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Analysis*, 7(1):73–108, 03 2012.

Merlise Clyde and Edward I George. Model uncertainty. *Statistical Science*, pages 81–94, 2004.

Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96:1348–1360, 2001.

Jianqing Fan and Jinchi Lv. A selective overview of variable selection in high dimensional feature space. *Statistica Sinica*, 20(1):101–148, 2010.

Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.

Richard Arnold Johnson, Dean W Wichern, et al. *Applied multivariate statistical analysis*, volume 4. Prentice hall Englewood Cliffs, NJ, 1992.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.

T. J. Mitchell and J. J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.

Naveen Naidu Narisetty and Xuming He. Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics*, 42(2):789–817, 2014.

Adrian E. Raftery, David Madigan, and Jennifer A. Hoeting. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92:179–191, 1998.

Veronika Ročková and Edward I George. EMVS: The EM approach to Bayesian variable selection. *Journal of the American Statistical Association*, 109(506):828–846, 2014.

Jun Shao and Xinwei Deng. Estimation in high-dimensional linear models with deterministic design matrices. *The Annals of Statistics*, 40(2):812–831, 2012.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1994.

Sijian Wang, Bin Nan, Saharon Rosset, and Ji Zhu. Random lasso. *The Annals of Applied Statistics*, 5(1):468–485, 2011.